arXiv:1704.01358v1 [cs.CV] 5 Apr 2017

Incremental Tube Construction for Human Action Detection

Harkirat S. Behl[1] Michael Sapienza[2] Gurkirt Singh[3] Suman Saha[3] Fabio Cuzzolin[3] Philip H. S. Torr[2]

[1] IIT Kanpur, [2] University of Oxford, [3] Oxford Brookes University

[1] harkirtb@iitk.ac.in, [2] {mikesapi,phst}@robots.ox.ac.uk,
[3] {gurkirt.singh-2015,suman.saha-2014,fabio.cuzzolin}@brookes.ac.uk

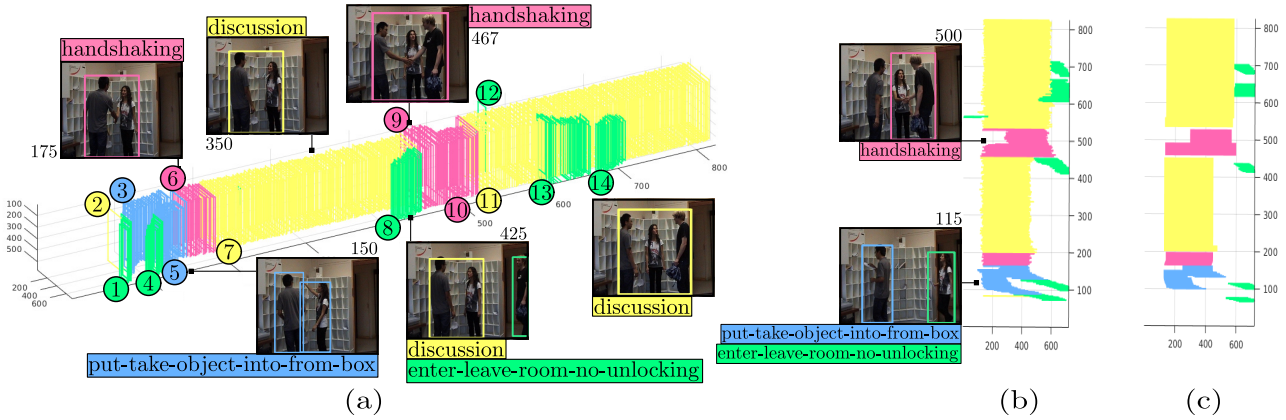

Figure 1: ***(a)*** *Our human action detection results on a video sequence from the LIRIS-HARL dataset [31]. The video starts when two people enter a room and put/take an object from a box (frame 150). They then shake hands (frame 175) and start having a discussion (see frame 350). In frame 450, another person enters the room, shakes hands, and then joins the discussion. Each action tube instance is numbered and coloured according to its action category. We selected this video to show that our tube construction algorithm can handle very complex situations in which multiple distinct action categories occur in sequence and at times concurrently. Our action detection algorithm is able to construct action tubes incrementally, and solves for the detection-window association and labelling with a single objective function.* ***(b)*** *The same video with action tubes drawn as viewed from above, compared to* ***(c)*** *the ground truth action tubes.*

Abstract

Current state-of-the-art action detection systems are tailored for offline batch-processing applications. However, for online applications like human-robot interaction, current systems fall short, either because they only detect one action per video, or because they assume that the entire video is available ahead of time. In this work, we introduce a real-time and online joint-labelling and association algorithm for action detection that can incrementally construct space-time action tubes on the most challenging action videos in which different action categories occur concurrently. In contrast to previous methods, we solve the detection-window association and action labelling problems jointly in a single pass. We demonstrate superior online association accuracy and speed (2.2ms per frame) as compared to the current state-of-the-art offline systems. We further demonstrate that the entire action detection pipeline can easily be made to work effectively in real-time using our action tube construction algorithm.

1. Introduction

Detecting human actions has been defined as the task of automatically predicting the start, end and spatial extent of various actions [12, 25, 31] by predicting sets of connected windows in time (called tubes) in which each action is enclosed, as illustrated in Fig. 1. Human action detection has gained huge popularity in the computer vision community due to its broad range of exciting applications. On the one hand, it is useful in 'offline' batch applications such as surveillance and the retrieval of video content in huge video collections. On the other hand, it can be used for online[1] human-robot interaction [24], an application in which both agents require instantaneous feedback and in which frame processing needs to be real-time and incremental.

Current state-of-the-art action detection methods have made remarkable progress on the aforementioned batch pro-

[1] By online we specifically mean that the algorithm is capable of processing its input frame-by-frame, without having the entire input available from the start.

cessing applications [8, 19, 23, 29], by dividing the action detection task into two steps: i) the extraction of independent per-frame action detection-windows, and ii) the linking (association) of the detection-windows to form action tubes. Due to the recent success of deep CNNs in object detection, frame-level action detections have improved drastically [23, 19]. The detection-window linking, however, still has several shortcomings.

Firstly, the linking is done by treating the video as a 3D block of frames (offline) [8, 19, 23, 29]. In this way, one can solve the problem globally. However, the number of possible solutions quickly explodes with the number of detection-windows, frames and action classes, so previous works [23] resort to picking the top detections in each frame.

Secondly, previous methods [23, 8, 19, 29] divide the problem into independent parts because that makes it easier to solve. In particular: i) The action tubes are constructed for each action class independently [23, 8, 29, 19], which can cause several action tubes with different action classes to overlap in the same space-time region (see Fig.2). They do not impose the constraint that only *one* action can take place in a region at an instant in time. They don't answer the question: "What action (singular) is being performed at a particular frame location?", thus ignoring the class assignment problem during association. Ignoring the class assignment during association also makes the algorithm slower and harder to scale to large number of action classes.

ii) The detection-window association and temporal localisation are performed as two separate optimisation steps [23, 29, 19]. This multi-part optimisation process is computationally redundant and expensive. Some effects of these shortcomings can be seen in Figs. 2 and 3.

Contributions. In this work, we tackle real-time human action detection applications and propose a novel incremental (single-pass) optimisation algorithm called OJLA (Online Joint Labelling and Association) that is able to construct and update action tubes as each new frame is added.

Our proposed action tube construction algorithm does away with multiple optimisation passes for detection-window association, temporal localisation and labelling [23]. Instead we formulate a single energy function and solve all of these tasks jointly and incrementally in a single pass. This implies that we do not perform action detection separately for each action class, thus avoiding the problem of detecting multiple co-located action tubes with different action classes (see Fig. 2). We demonstrate that our action tube construction algorithm outperforms previous batch methods [8, 19, 23, 29] in terms of accuracy and speed (detection-window association takes ∼2.2ms per frame or ∼450 fps when averaged over all videos on UCF101), by virtue of our simplified single-pass optimisation coupled with an appropriate energy function.

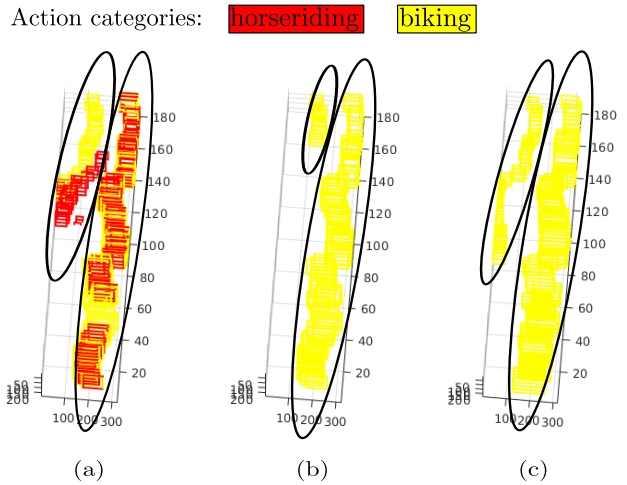

Figure 2. **(a)** The quality of action tubes generated by the algorithm of Saha *et al*. [23], compared to **(b)** the action tubes generated by our method with an improved cost function. **(c)** The ground truth action tubes.

To further show that our action tube construction algorithm can find application in real-time action detection applications like human-robot interaction, we demonstrate a real-time action detection system (§ 5), which combines the real-time frame-level window detector of [17] and our OJLA algorithm, and which outperforms current state-of-the-art offline methods.

Paper organization. The rest of the paper is organised as follows: Section 2 presents the related work on action detection and association methods. Our OJLA association method is detailed in section 3, and followed by a brief description on frame-level detection-window generation (§ 4). The experiments are discussed in section 5.

2. Related Work

Multiple action detection in space and time. Several works attempt to localise actions in space alone using a single-frame detector [3, 32], or in time alone [4, 6]. Here we focus on space-time action detection, defined by a set of continuous sequences of 2D detection-windows in time to form action tubes.

Laptev *et al*. [15] pioneered one of the first works on action detection in realistic scenarios, with a boosted 3D space-time window classifier combined with a single frame action detector. An alternative approach by Willems *et al*. extended exemplar based object detection to action detection [30]. In order to reduce the action search space, Klaser *et al*. [14] first generated action hypothesis by detecting and tracking humans, and subsequently classified the space-time 'tracks'. Frame-by frame detections were generated by

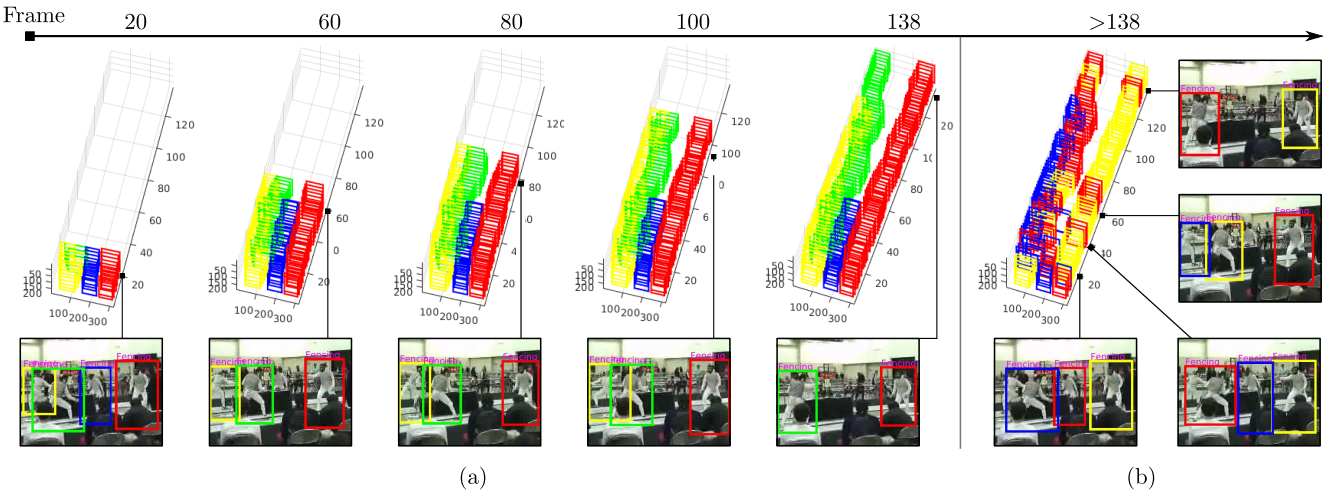

Figure 3. An illustration of the difference between **(a)** our online, incremental algorithm for constructing action tubes, and **(b)** one of the state-of-the-art offline methods by [23], using the exact same scores coming from a two-stream CNN. The video from UCF-101 contains four humans competing in a fencing tournament. Each colour indicates a different action tube present in the video. Notice that in **(a)**, at each new frame in time, our method is able to predict the detection-window association and class of each currently active action tube. Solving for the association and class assignment in one pass allows us to reduce the window switching effects **(b)** observed in the results by prior art [23].

training an upper body HoG detector, and associating the detected windows in time using KLT-tracked features.

More recent works for action localisation and classification based on Dense Trajectory features [26, 11, 34] developed space-time and class independent proposal algorithms (offline). Note that Yu *et al*. [34] add a constraint to avoid generating redundant overlapping proposals. This is in agreement with our observation that only a single action may happen at one space-time location.

Since the advent of deep learning, the detection of action hypothesis has improved dramatically. We take inspiration from [8, 29, 19, 23], and use the detections generated by fine-tuning a CNN on the action categories present in a particular dataset. In fact we use the exact same detections generated by the state-of-the-art method by Saha *et al*. [23] in order to be able to make a fair comparison, and in order to focus the merits of our work on the construction of action tubes.

Association and labelling. Some research has been dedicated to action localisation in the image plane, with no association over time [18]. Other works [19, 8, 29] provide an evaluation on a per-frame basis. Without association, a robot would not be able to recognise that a walking action over multiple frames is in fact a single action[2]. Association in this context involves the assignment of detection-windows to existing action tubes. Most of the current state-of-the-art methods in action detection [8, 29, 19, 23] do link the detection-windows to generate space-time action tubes, however, so far this has been an offline procedure. For instance [8, 19, 23] have used the Viterbi algorithm [27] iteratively.

In the tracking community however, online data association methods are widely used. These include the multi-hypothesis tracker (MHT) [20, 2] and the Joint probabilistic data association (JPDA) algorithm [5, 9]. A more computationally efficient algorithm extending JPDA was proposed by Hamid *et al*. [9] in which the 'm-best' solutions were calculated instead of all the solutions. By design, the mentioned multi-object tracking algorithms [20, 9] can only track multiple objects of a single class, such as 'person' or 'biological cell', whereas in human action detection, tubes with multiple classes (action categories) are present over the course of a video (see Fig.1). Thus we bridge the gap between human action detection [19, 23] and incremental multiple-object tracking [5]. In doing so, we propose a joint labelling and association algorithm inspired by [9] which can handle multiple co-occurring objects and classes over time.

Incremental, online video stream processing. Despite recent advances in online action recognition [33, 35], current systems are only tailored for the recognition of a single action category per video frame. For example, Yeffet *et al*. [33] proposed an online action recognition system in which grid based video regions were represented as a histogram of Trinary strings. Whilst this approach is fast and online, it only encodes motion information, and crucially is limited

[2]Note that it is not easy for humans to determine whether an action detection system is connecting actions in time from a video displaying bounding box predictions, since we perform the data association effortlessly whilst watching.

to recognise a single action class per frame.

Yu *et al*. [35] designed a very fast action recognition system based on ensembles of random decision trees that quickly translate interest points into visual codewords for classification. The downside is that it also only allows single action classification per group of frames. We improve over previous efforts in online and incremental action recognition [33, 13] by allowing multiple simultaneous actions to be detected concurrently.

3. Methodology

Inspired by previous methods who split the action detection task into two steps, we also first extract action detection-windows in each video frame independently (§ 4), and then perform linking of the detection-windows to construct action tubes, as detailed in what follows. A frame level action detection-window is a spatial bounding box enclosing a particular action category $l \in \mathcal{L}$, where $\mathcal{L}$ is the set of action categories present in a particular dataset; and an action tube represents a set of continuous detection-windows (without breaks) in time, which share the same action category $l \in \mathcal{L}$ [19, 23].

3.1. Online joint labelling and association (OJLA)

Problem formulation. At any time t, we seek to find the association between a set of currently active tubes $\{x^i_{t-1}\}$, $i \in \mathcal{N} = (1, \ldots, N)$, and a new set of detection-windows $\{y^j_t\}$, $j \in \mathcal{M}_0 = (0, 1, \ldots, M)$ from the current time-step, and also update the labels of the current tubes. Here '0' is a placeholder for a dummy (or missed) detection. This is a one-to-one matching problem in which each tube should have a unique edge connecting it to a detection-window.

The best assignment can be formulated as the solution to the problem of finding those connections which minimize an appropriate cost function f:

$$\mathbf{a}^* = \arg\min_{\mathbf{a}\in\tau} f(\mathbf{a}), \tag{1}$$

where $\mathbf{a}$ is a binary assignment vector in which a^j_i signifies that tube i is connected to detection-window j, and τ is the space of all valid combinations of detection-to-tube assignments. Each valid assignment should satisfy the following constraints: i) each tube is uniquely assigned to a single detection (2a), and ii) each detection (except for dummy hypothesis $j = 0$) is assigned to at most one tube (2b). Thus, τ is defined as the following set of binary vectors:

$$\tau = \Big\{ \mathbf{a} = (a^j_i)_{i\in\mathcal{N}, j\in\mathcal{M}_0} \,|\, a^j_i \in \{0, 1\}, \tag{2}$$

such that

$$\sum_{j\in\mathcal{M}_0} a^j_i = 1 \qquad \forall i \in \mathcal{N} \tag{2a}$$

$$\sum_{i\in\mathcal{N}} a^j_i \leq 1 \qquad \forall j \in \mathcal{M} \Big\}, \tag{2b}$$

where $a^j_i = 1$ indicates that tube i is matched to detection j and $a^j_i = 0$ indicates that tube i is not matched to detection j. Note that $\mathcal{M}$ is the set of detection-window indices excluding the zero dummy index. Thus, $\mathbf{a} \in \tau \subseteq B^{N\times(M+1)}$ is a binary vector which represents one valid solution of the one-to-one matching problem.

Cost formulation. We denote the total cost for an assignment as

$$f(\mathbf{a}) = \mathbf{c}^\top \mathbf{a}, \tag{3}$$

where $\mathbf{c}$ is a cost vector in which c^j_i is the cost for associating detection index $j \in \mathcal{M}_0$ with tube $i \in \mathcal{N}$ at time t. Finding the best assignment $\mathbf{a}^*$, from Eqn. 1, Eqn.2 and Eqn. 3, thus becomes a binary linear program.

For convenience, we define the cost c^j_i in terms of a similarity score by inverting it:

$$c^j_i = \begin{cases} c_0, & \text{if } j = 0. \\ 1/s^j_i, & \text{otherwise,} \end{cases} \tag{4}$$

where c_0 is a constant for assigning the tube to no-detection, and s^j_i is the score defining the similarity between tube i and detection-window j.

A tube's state at time $t-1$ is a composition of the bounding box coordinates, action scores, and it's action label:

$$x^i_{t-1} = (\mathbf{b}^i_{t-1}, \mathbf{z}^i_{t-1}, l^{i*}), \tag{5}$$

where $\mathbf{b}$ represents the detection-window parameters, $\mathbf{z}$ is a vector of action scores per category, and l^* is the assigned action label. The detection-windows at time t are composed only of boxes and scores: $y^j_t = (\mathbf{b}^j_t, \mathbf{z}^j_t)$.

The similarity score is defined in terms of a labelling score function ψ and an overlap score function ψ_o:

$$s^j_i = \psi(x^i, y^j) + \lambda_o \psi_o(x^i, y^j), \tag{6}$$

where λ_o is a scalar that controls the relative weighting between the two terms, and the overlap is calculated as the intersection-over-union between the tube and the detection-window.

The labelling score function is calculated as the maximum sum of unary tube and detection-window scores, minus a Potts penalty to encourage label smoothness, over all possible action labels:

$$\psi(x^i, y^j) = \max_{l\in\mathcal{L}} \Big(\mathbf{z}^i_{t-1}(l) + \mathbf{z}^j_t(l) - \bar{\psi}(l^{i*}, l) \Big). \tag{7}$$

Here $\mathbf{z}^i_{t-1}(l)$ is the unary score for the tube taking action label l, and $\mathbf{z}^j_t(l)$ is the corresponding score for the detection-window at time t. The label l^{i*} is the assigned action label

of the tube i, and $\bar{\psi}(l', l)$ is the Potts penalty incurred by switching from action l' to l, where

$$\bar{\psi}(l', l) = \begin{cases} 0 & \text{if } l' = l \\ \alpha & \text{otherwise.} \end{cases} \quad (8)$$

The labelling term is motivated by the fact that we not only want to estimate the detection-window corresponding to the same tube, but also the action label which is being performed. Thus, Eqn. (7) decides the most probable action category assuming that detection j forms part of tube i, taking care of the labelling whilst computing the score. Since the labelling score function makes action category decisions on each detection window, it will take care of 'temporally trimming' of the action tube if a no-action category is present in the label set. In computing Eqn. (7), we found in practice that taking the sum of action scores from the tube over a time window of n frames improves robustness to spurious predictions.

Our cost formulation differs significantly from [19, 23, 29] as we do the association, labelling, smoothing and temporal trimming in a single pass.

Optimisation. The optimal assignment $\mathbf{a}^*$ may be found using an off-the-shelf optimisation solver for binary linear programs. Instead of picking the best solution, [10] showed the improvements that can be obtained by marginalising over several (m-best) solutions, as described in the following section.

3.2. Marginalizing over m-best

In [10, 1], the authors argue that in many cases it is beneficial to look at the m-best solutions of the one-to-one matching problem instead of just using the best solution. A brief motivation is that we are using a relatively simple cost function made up of unary and pairwise terms to model a complicated real-world application such as action detection. In practice picking the best solution may not always give good results, as there may be numerous competing solutions which are almost equally likely.

We therefore marginalize $f(\mathbf{a})$ over the m-best solutions of the matching space τ. Let $\tau_i^j \subset \tau$ be the subset of τ which includes all the solutions in which tube i is matched to detection j, the marginalized cost q_i^j for assigning detection index $j \in \mathcal{M}_0$ to tube $i \in \mathcal{N}$ is calculated as:

$$q_i^j = -\log \sum_{\left\{\mathbf{a}_\mathbf{k}^* | \forall k \in [m], a_i^j = 1\right\}} e^{-f(\mathbf{a}_\mathbf{k}^*)}, \quad (9)$$

where $\mathbf{a}_k^*$ is the k^{th} best solution to Eqn. (1), as in [9]. The k^{th} best solution is calculated using the binary tree partitioning (BTP) algorithm [9]. BTP removes the redundant constraints and computes the objective as a series of second-best solutions.

3.3. Tube-state update

The marginalised costs $(q_i^j)_{j \in \mathcal{M}_0}$ are normalised for each tube, and the detection corresponding to the minimum cost is assigned to tube i. Let j^* be the index of the detection matched to tube i, and l_t^* be the predicted action label calculated from Eqn. 7 for edge (i, j^*). If the predicted action label is the same the label of the action tube, the tube's state is then updated as:

$$x_t^i = (\mathbf{b}_t^{j^*}, \mathbf{z}_t^{j^*}, l^{i*}). \quad (10)$$

If l_t^* is not the same as l^{i*}, the action tube is terminated, and a new tube with label l_t^* is initiated.

Initiation and termination. In order to initiate and terminate action tubes, we took inspiration from a heuristic algorithm proposed for multi-target tracking [22], and adapted it to our action detection application.

Firstly, any detection-window that is not claimed by an existing action tube, but that has a high score for a particular action (excluding no-action) is initiated as a new tube. Secondly, a tube is terminated if its number of consecutive missed detection assignments reach a specific threshold.

4. Frame Level Action Detections

4.1. Detection

In order to generate frame-level detection-windows for a video, we use two parallel streams of Region Proposal Networks [21], which have been fine-tuned by [23] for predicting action or no-action, instead of object or no-object. The first network takes an RGB image as input, whereas the second network takes the optical flow calculated between two adjacent video frames. These action proposals are then fed through the Fast R-CNN network of [7]. For each region proposal, the Fast R-CNN network outputs a set of regressed bounding boxes per category and corresponding scores. Throughout this work, we use the bounding boxes and scores that we generated by running the publicly available code of [23][3].

4.2. Fusion of appearance and motion cues

Both the appearance and motion CNN networks output category-specific detection-windows and scores for each region proposal. We map the set of boxes and scores that originated from one region proposal back to a single box with a vector of scores by performing a weighted average of the boxes, where the weight corresponds to the action-specific score. The next step involves fusing detection-windows between the two independent appearance and motion CNNs, we use a similar approach as Saha et al. [23] with the main

[3] http://sahasuman.bitbucket.org/bmvc2016

difference being that we perform the fusion for all the categories directly.

5. Experiments and Results

Improving computational performance. In order to improve computational performance, we make use of a gating procedure to eliminate unlikely detection-to-tube edges. In practice, we exclude the edges whose corresponding Euclidean distance between the tube and detection centers is larger than a threshold d. This translates into exploring only the local region around the tube's position in the previous frame. Because of this, our algorithm is able to use one order of magnitude more detection-windows for association than [23], who use the top 5 windows per frame.

We fix the maximum number of m-best solutions (§ 3.2) to 100 for all our experiments.We found that this approximation gives as much accuracy as marginalizing over all the solutions, but is ~ 10 times faster.

Datasets. We perform an extensive evaluation of our approach on three of the most challenging datasets on action detection, namely, UCF101 [25], J-HMDB-21 [12] and LIRIS-HARL [31]. The J-HMDB-21 dataset has 21 action categories and each video only contains a single action which extends to the entire length of the video. Providing more difficulty, the UCF-101 dataset has 24 action categories and contains sequences in which multiple action tubes of the same category are present simultaneously with varying temporal extent. Nevertheless, only one action category is present in each video. Our primary focus though is the LIRIS-HARL dataset, in which multiple action tubes, and action categories may be present simultaneously. The qualitative results of our algorithm on videos in which different action categories occur concurrently are shown in Figs. 1, 6 and 7. This provides a more realistic benchmark with which to evaluate our system.

Performance measures. We report the mean average precision at various intersection overlap thresholds in order to compare to the state-of-the-art. Whereas Saha *et al.* [23] reported results for overlap thresholds from 0.1 to 0.7 in increments of 0.1, we feel that thresholds below 0.5 are too low to be of any practical use. Instead we report the results for thresholds of 0.5 and above, and also report the main metric used to benchmark the object detection and instance segmentation methods in the MS-COCO [16] challenge: the mAP averaged over thresholds from 0.5 to 0.95 in steps of 0.05. The quantitative results for our experiments are detailed in Tables 1, 2 and 3. In Tables 2 and 3, we also report the results for thresholds of 0.3 and 0.4 respectively, for a fair comparison with [29, 19, 23].

Table 1. Quantitative action detection results on J-HMDB-21.

mAP @ space-time overlap threshold δ	.5	.6	.7	.5:.95
ActionTube [8]	53.3	–	–	–
Wang *et al.* [28]	56.4	–	–	–
STMH [29]	60.7	–	–	–
MR-TS R-CNN [19]	**73.10**	–	–	–
Saha *et al.* [23]	71.50	**68.73**	**56.57**	–
Ours (with OJLA)	60.27	57.82	49.77	34.59
Ours (with OJLA multilabel)	67.29	61.48	49.09	**36.08**

Table 2. Quantitative action detection results UCF-101.

mAP @ space-time overlap threshold δ	.3	.5	.6	.5:.95
STMH [29]†	37.82	–	–	–
MR-TS R-CNN [19]	**65.78**	32.07	11.25	07.26
Saha *et al.* [23]	55.46	36.37	26.87	14.37
Ours (with OJLA)	48.58	31.13	24.70	12.67
Ours (with OJLA multilabel)	56.87	**37.13**	**29.64**	**15.09**

† This method uses different annotations from others

Table 3. Quantitative action detection results on LIRIS-HARL.

mAP @ space-time overlap threshold δ	.4	.5	.6	.5:.95
Saha *et al.* [23]	28.03	21.36	–	–
Ours (with OJLA)	**29.17**	**21.77**	**9.59**	**6.16**
Ours (with OJLA multilabel)	26.42	16.93	4.43	4.30

OJLA with multiple labels for comparison. In section 3 we outlined our original multiclass algorithm, OJLA. With a small modification, our algorithm can also be made to predict several action tubes with different category labels in the same space-time region. In contrast to the multiclass approach, where only the tube corresponding to the action with maximum score in constructed, in the multilabel approach, we construct tubes for each action independently. A Tube for label l is constructed if the score in Eqn. 7 corresponding to action category l is larger than a threshold value. This is the only difference between the quantitative results 'Ours (with OJLA)' and 'Ours (with OJLA multilabel)' listed in Tables 1, 2 and 3.

Results. We compare the quantitative results that we achieved to the current state-of-the-art in Tables 1, 2 and 3. Despite using the same detection-windows and action scores as used in [23], and the fact that our method is incremental, we achieve state-of-the-art results on the more difficult UCF-101 and LIRIS-HARL benchmarks. Moreover our algorithm is approximate given that we can only look at the present and past frames, whereas offline approaches can make use of the entire information carried by the test video. The fact that our online greedy method is better than current offline methods for action detection suggests that there is scope for improving offline action tube construction.

On J-HMDB-21, only one action category and tube are present in each video, which is why it is the easiest of the three datasets that we evaluate on. Only on J-HMDB-21

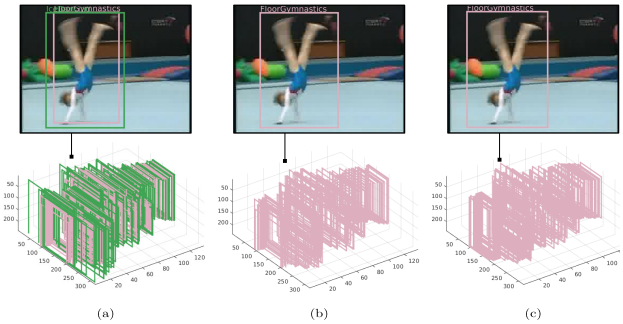

Figure 4. **(a)** Action tubes predicted by the method of [23] with 'ice dancing' (green) and 'gymnastics' (pink) labels in the same space-time region. **(b)** In contrast our method assigns a single action label to a space-time region, eliminating the possibility of predicting multiple co-located action tubes. **(c)** The ground truth action tube from the UCF-101 dataset.

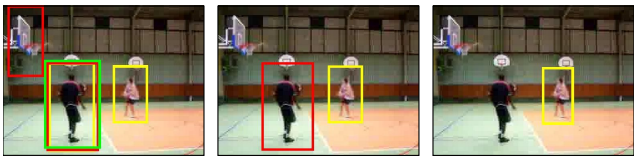

Figure 5. **(a)** Another example in which [23] predicts multiple co-located action tubes. This is a basketball sequence form UCF-101; the colours symbolise the following action categories: (green - Soccer Juggling), (red - Cricket Bowling), (yellow - Basketball). **(b)** Our multiclass predictions. **(c)** The ground truth annotation.

competing offline methods do better than ours; note that this dataset does not present the sort of problematic situations our algorithm is designed to deal with.

On UCF-101, our online multilabel method surpasses the state-of-the-art across all detection thresholds. Example qualitative results from UCF-101 are shown in Figs. 4 and 5.

It is noteworthy that on LIRIS-HARL, a dataset which includes several action categories, our OJLA algorithm can truly leverage on its unique properties to get the best results across the board. Qualitative results from LIRIS-HARL are shown in Figs. 6 and 7.

It is surprising that our online method turns out to be better than the reported offline methods. This is partly due to the fact that we are solving the association and labelling jointly, instead of using multiple optimisation passes [23]. This single pass also allows us to keep track of humans switching between several action categories (see Figs. 6 and 7), in contrast to previous methods [8, 19, 23, 29]. Another factor includes our ability to look at an order of magnitude more detection-windows per frame as compared to [23].

Discussion. One of the main differences between our proposed method and those of previous methods is that we treat each region in space as belonging to one action category at

Figure 6. Qualitative results of our method on the LIRIS-HARL dataset. First (top), a woman walks into a room whilst a man stands in front of a whiteboard. The two people then 'shake hands' and start a 'discussion'. Notice how our algorithm is able to handle situations in which multiple actions occur concurrently and/or sequentially. Next (middle) a person 'enters/leaves a room without unlocking', then 'puts-takes an object from a box', and again 'enters/leaves a room without unlocking'. Finally (bottom) a man holds a 'telephone conversation'; again the system mislabels the beginning and end of the action by detection a 'put/take object into/from box' action immediately preceding and following the 'telephone conversation'.

one instant in time, a multiclass approach, in contrast to a multilabel approach in which each region in space may assume multiple action labels at one point in time (a method which we argue is flawed for this action detection application). The difference this basic assumption makes in terms of qualitative results can be seen in Figs. 2, 4 and 5. For instance in Fig.4a, the method of [23] predicts that the girl performing a gymnastics action is also 'ice dancing', an impossible combination. The different approaches are asking subtly different questions, the multilabel approaches are asking: 'Is this action happening or not?' (independently for each action category), whilst our multiclass approach is asking 'Which action is happening now?' (including no action). Since our method is deciding on the best possible action category for each tube at each frame, we do not run the risk of labelling each space-time region with multiple action categories.

Another example of the difference between these two approaches can be seen in Fig. 5, where two people are playing basketball. In this sequence, the person on the left is waiting to receive a ball and is not performing any action. However the multilabel approach of [23] labels the person with three actions; in many cases the fact that people are labelled with several actions can help to boost the results if at least one of

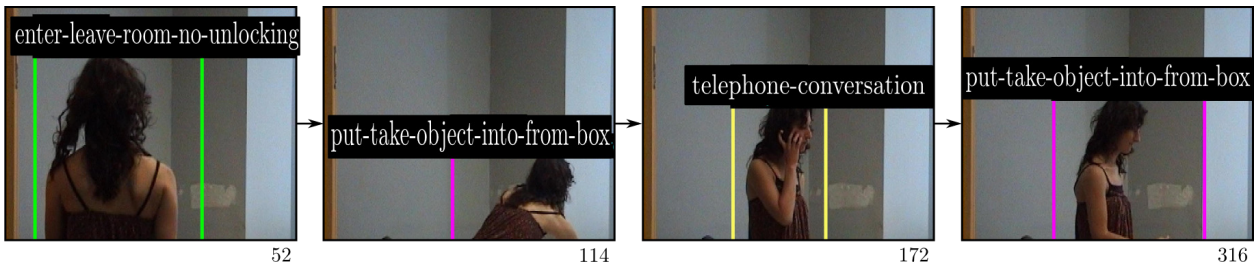

Figure 7. Video frames from video 27 in the LIRIS-HARL dataset showing a single person performing several actions. From left, the person first 'enters the room' (frame 52), then has a 'telephone conversation' (frame 172). Before and after the telephone conversation, the human is detected as 'putting/taking an object from a box'. The full video sequence is shown in the supplementary video.

the categories is correct (see Table 2). To prove the point, we implement a similar multilabel approach 'Ours (OJLA multilabel)' for comparison.

Even though one would expect the multilabel quantitative results to become worse (see the qualitative comparison in Fig. 4), they in fact increase the mAP performance, as shown in Table 2. It is noteworthy that the two approaches (multiclass vs multilabel) being compared use the exact same association algorithm (OJLA). However, the multilabel output predicts overlapping tubes with different action labels whereas our original multiclass algorithm predicts the exact same tube with just a single label. The reason for this discrepancy in results may be interpreted as follows. Consider a situation in which the detection-window association is accurate, but the labelling is poor. In this case a multilabel approach helps to improve results because predicting multiple tubes with different labels will increase the chances of getting a true positive detection. The multiclass approach, however, will suffer if the association is accurate but the labelling is poor, as its incorrect labelling (greedy) will not result in a true positive. We observed this behaviour on the J-HMDB-21 and UCF-101 datasets, where the mAP are quite high due to good detection-window scores and associations. When the association is less good, as observed on LIRIS-HARL, then predicting multiple overlapping tubes (as in the multilabel approach) won't help because neither of the tubes will have a correct label and all predictions would be counted as false positives. This effect was observed on the LIRIS-HARL dataset (see Table 3). Thus, we can observe that predicting multiple overlapping tubes can improve mAP results in relatively easy situations, however these improvements will not transfer to more complex scenarios. Moreover, it seems as if the multilabel approaches are flawed since they do not solve the general problem of action detection in which only one action can take place in a given place and time, something which current performance metrics do not take into consideration.

In real-life, more than one action happens in a video, and for that reason we think that the LIRIS-HARL dataset [31] is more appropriate for evaluating human action detection algorithms. An example of our results on a complex video from LIRIS-HARL can be seen in Fig. 1. Note that, we are the first to show detailed qualitative results on action videos with multiple distinct actions happening simultaneously.

Table 4. Quantitative action detection results UCF-101.

mAP @ space-time overlap threshold δ	.3	.5	.6	.5:.95
MR-TS R-CNN [19]†	**65.78**	32.07	11.25	07.26
Saha *et al.* [23] †	55.46	36.37	26.87	14.37
Ours (SSD with OJLA) *	46.60	30.12	24.99	13.96
Ours (SSD with OJLA multilabel) *	58.60	**39.16**	**32.69**	**17.94**

* Real-time † Offline

Real-time action detection. Here we demonstrate that our OJLA tube construction algorithm can be used to create a real-time action detection system by pairing it with the Single Shot Detector (SSD) [17] for frame-level action detections. We train the SSD network[4] [17] on the RGB images in UCF-101 using the same network architecture and training procedure as used in [17]. The SSD network takes ∼21.7ms per frame and OJLA takes ∼2.2ms per frame. The speed of the overall detection system is thus limited by the speed of the frame-level detection network, which is ∼46fps. The quantitative results of our real-time action detection system are reported in Table. 4. The relative increase in performance as compared to Table. 2 further demonstrates that the overall results can be improved with more accurate frame-level detections.

6. Conclusion

In this work we have proposed an incremental and online algorithm for the construction of space-time action tubes. Interestingly, we have shown that our online and method achieves state-of-the-art results in action detection despite competing with offline methods that have the potential to find globally optimum solutions. We have shown that by constructing a suitable energy function, we are able to solve the action detection-window association and labelling jointly in a single pass.

[4] https://github.com/weiliu89/caffe/tree/ssd

References

[1] D. Batra. An efficient message-passing algorithm for the m-best map problem. *arXiv preprint arXiv:1210.4841*, 2012. 5

[2] I. J. Cox and S. L. Hingorani. An efficient implementation of reid's multiple hypothesis tracking algorithm and its evaluation for the purpose of visual tracking. *IEEE Transactions on pattern analysis and machine intelligence*, 18(2):138–150, 1996. 3

[3] C. Desai and D. Ramanan. Detecting actions, poses, and objects with relational phraselets. In *Proc. European Conf. Computer Vision*, pages –, 2012. 2

[4] O. Duchenne, I. Laptev, J. Sivic, F. Bach, and J. Ponce. Automatic annotation of human actions in video. In *Proc. Int. Conf. Computer Vision*, 2009. 2

[5] T. Fortmann, Y. Bar-Shalom, and M. Scheffe. Sonar tracking of multiple targets using joint probabilistic data association. *IEEE Journal of Oceanic Engineering*, 8(3):173–184, Jul 1983. 3

[6] A. Gaidon, Z. Harchaoui, and C. Schmid. Actom sequence models for efficient action detection. In *IEEE Int. Conf. on Computer Vision and Pattern Recognition*, 2011. 2

[7] R. Girshick. Fast R-CNN. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1440–1448, 2015. 5

[8] G. Gkioxari and J. Malik. Finding action tubes. In *IEEE Int. Conf. on Computer Vision and Pattern Recognition*, 2015. 2, 3, 6, 7

[9] S. Hamid Rezatofighi, A. Milan, Z. Zhang, Q. Shi, A. Dick, and I. Reid. Joint probabilistic data association revisited. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 3047–3055, 2015. 3, 5

[10] S. Hamid Rezatofighi, A. Milan, Z. Zhang, Q. Shi, A. Dick, and I. Reid. Joint probabilistic matching using m-best solutions. In *IEEE Int. Conf. on Computer Vision and Pattern Recognition*, June 2016. 5

[11] M. Jain, J. Van Gemert, H. Jégou, P. Bouthemy, and C. G. Snoek. Action localization with tubelets from motion. In *IEEE Int. Conf. on Computer Vision and Pattern Recognition*, pages 740–747, 2014. 3

[12] H. Jhuang, J. Gall, S. Zuffi, C. Schmid, and M. Black. Towards understanding action recognition. *Proc. Int. Conf. Computer Vision*, 2013. 1, 6

[13] Y. Ke, R. Sukthandar, and M. Hebert. Volumetric features for video event detection. *Int. Journal of Computer Vision*, 88(3):339–362, 2010. 4

[14] A. Kläser, M. Marszałek, C. Schmid, and A. Zisserman. Human focused action localization in video. In *International Workshop on Sign, Gesture, Activity*, 2010. 2

[15] I. Laptev and P. Pérez. Retrieving actions in movies. In *Proc. Int. Conf. Computer Vision*, 2007. 2

[16] T. Lin, M. Maire, S. J. Belongie, L. D. Bourdev, R. B. Girshick, J. Hays, P. Perona, D. Ramanan, P. Dollár, and C. L. Zitnick. Microsoft COCO: common objects in context. In *Proc. European Conf. Computer Vision*, 2014. 6

[17] W. Liu, D. Anguelov, D. Erhan, C. Szegedy, S. Reed, C.-Y. Fu, and A. C. Berg. SSD: Single shot multibox detector. *arXiv preprint arXiv:1512.02325*, 2015. 2, 8

[18] K. Mikolajczyk and H. Uemura. Action recognition with motion-appearance vocabulary forest. In *IEEE Int. Conf. on Computer Vision and Pattern Recognition*, 2008. 3

[19] X. Peng and C. Schmid. Multi-region two-stream R-CNN for action detection. In *ECCV 2016 - European Conference on Computer Vision*, Amsterdam, Netherlands, Oct. 2016. 2, 3, 4, 5, 6, 7, 8

[20] D. Reid. An algorithm for tracking multiple targets. *IEEE transactions on Automatic Control*, 24(6):843–854, 1979. 3

[21] S. Ren, K. He, R. Girshick, and J. Sun. Faster R-CNN: Towards real-time object detection with region proposal networks. In *Advances in Neural Information Processing Systems*, pages 91–99, 2015. 5

[22] S. H. Rezatofighi, S. Gould, R. Hartley, K. Mele, and W. E. Hughes. Application of the imm-jpda filter to multiple target tracking in total internal reflection fluorescence microscopy images. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 357–364. Springer, 2012. 5

[23] S. Saha, G. Singh, M. Sapienza, P. H. Torr, and F. Cuzzolin. Deep learning for detecting multiple space-time action tubes in videos. In *Proc. British Machine Vision Conference*, 2016. 2, 3, 4, 5, 6, 7, 8

[24] Y. Sakagami, R. Watanabe, C. Aoyama, S. Matsunaga, N. Higaki, and K. Fujimura. The intelligent ASIMO: system overview and integration. In *IEEE/RSJ Int. Conf. on Intelligent Robots and Systems*, 2002. 1

[25] K. Soomro, A. R. Zamir, and M. Shah. UCF101: A dataset of 101 human action classes from videos in the wild. Technical report, CRCV-TR-12-01, 2012. 1, 6

[26] J. C. van Gemert, M. Jain, E. Gati, and C. G. Snoek. APT: Action localization proposals from dense trajectories. In *Proc. British Machine Vision Conference*, volume 2, page 4, 2015. 3

[27] A. Viterbi. Error bounds for convolutional codes and an asymptotically optimum decoding algorithm. *IEEE Trans. Inf. Theor.*, 13(2):260–269, Sept. 2006. 3

[28] L. Wang, Y. Qiao, X. Tang, and L. V. Gool. Actionness estimation using hybrid fully convolutional networks. In *CVPR*, pages 2708–2717, 2016. 6

[29] P. Weinzaepfel, Z. Harchaoui, and C. Schmid. Learning to track for spatio-temporal action localization. In *IEEE Int. Conf. on Computer Vision and Pattern Recognition*, June 2015. 2, 3, 5, 6, 7

[30] G. Willems, J. H. Becker, T. Tuytelaars, and L. V. Gool. Exemplar-based action recognition in video. In *Proc. British Machine Vision Conference*, 2009. 2

[31] C. Wolf, J. Mille, E. Lombardi, O. Celiktutan, M. Jiu, M. Baccouche, E. Dellandra, C.-E. Bichot, C. Garcia, and B. Sankur. The LIRIS Human activities dataset and the ICPR 2012 human activities recognition and localization competition. Technical Report RR-LIRIS-2012-004, LIRIS UMR 5205 CNRS/INSA de Lyon/Université Claude Bernard Lyon 1/Université Lumière Lyon 2/École Centrale de Lyon, Mar. 2012. 1, 6, 8

[32] B. Yao and L. Fei-Fei. Action recognition with exemplar based 2.5d graph matching. In *Proc. European Conf. Computer Vision*, 2012. 2

[33] L. Yeffet and L. Wolf. Local trinary patterns for human action recognition. In *Proc. Int. Conf. Computer Vision*, 2009. 3, 4

[34] G. Yu and J. Yuan. Fast action proposals for human action detection and search. In *IEEE Int. Conf. on Computer Vision and Pattern Recognition*, pages 1302–1311, 2015. 3

[35] T.-H. Yu, T.-K. Kim, and R. Cipolla. Real-time action recognition by spatiotemporal semantic and structural forests. In *Proc. British Machine Vision Conference*, 2010. 3, 4